Optimizing facility operation in high density data center environments

technology brief

Abstract ... 3

Introduction ... 3

Assessing facility requirements ... 3
- Using sizing tools to understand server and facility requirements ... 4
- Data center assessment ... 5

Component power and cooling ... 5
- Processor power and cooling ... 5
  - Processor P-state ... 5
  - P-state management ... 6
- Efficient practices for servers and enclosures ... 6
  - Understanding server power utilization and heat generation ... 6
  - Greater density with multi-core processors ... 7
  - Blade enclosures with integrated cooling hardware ... 7
  - Onboard thermal sensors and temperature management ... 7
  - Pooled power ... 7
  - High-line power efficiency ... 7
  - Power management and reporting of individual servers ... 8
  - Server virtualization using virtual machine technology ... 9
- Efficient practices at the rack level ... 9
  - Rack configuration tools ... 9
  - Methods for calculating power supply requirements ... 10
  - Power protection and management ... 10
  - Rack based power distribution ... 10
  - High-efficiency power supplies ... 10
  - Understanding Internal airflow ... 10
  - Liquid cooled racks ... 11

Efficient practices for facility-level power and cooling ... 12
- Power management and reporting ... 13
  - Comprehensive management applications ... 13
  - Power Capping ... 13
- Airflow distribution for high-density data centers ... 13
  - Raised floors ... 15
  - Air supply plenum ... 15
  - Ceiling return air plenum ... 17
  - Dual supply air plenums ... 17
  - Perforated tiles ... 17
  - Rack geometry ... 18
  - Cooling footprint ... 18
  - Hot and cold aisle spacing ... 18
  - Row configuration ... 18

Downloaded from www.Manualslib.com manuals search engine

Closely coupled cooling 19
Computer room oir conditioners 19
Copocity of CRAC units 19
Plocement of CRAC units 20
Dischorge velocity 21
Advonced thermol monogement techniques 22
Thermol Assessment Services 22
Dynomic Smort Cooling 24
Summory 25
For more informotion 26
Coll to oction 26

Downloaded from www.Manualslib.com manuals search engine

Abstract

This paper describes issues of high processor and server density within existing data center infrastructures. It identifies methods to optimize the effectiveness of power and cooling resources in facilities that are deploying high-density equipment or that are already fully populated with high-density equipment. The intended audience for this paper includes IT managers, IT administrators, facility planners, and operations staff. Managers considering strategic planning activity such as facility upgrades or new facility planning should consult the "Data Center Cooling Strategies" technical brief available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01153741/c01153741.pdf.

Introduction

Existing data centers have had three major operational and fiscal constraints – power, cooling and space. As existing data centers are required to support increasingly dense configurations, power and cooling requirements can outstrip the capabilities of the data center infrastructures. In fact, the issue of space becomes moot because existing data centers are projected to run out of power before they run out of space.

HP understands that individual data center challenges like rack and processor level power and cooling cannot be viewed and managed as disconnected issues. HP has developed solutions to the fundamental issues of power and cooling at the processor, server, rack, and facility infrastructure level. HP also has developed proven management tools to provide a unified approach to managing power and cooling in data centers. This document describes issues, tools, and solutions for optimizing power and cooling at each level of the data center: processor, server, rack, and facility.

The first section of this paper describes facility assessment and how the component and server level issues affect power and cooling, while the last half of the paper describes rack-level, facility layout, and data center level approaches.

Assessing facility requirements

Data center facilities face unprecedented power and cooling requirements. For example, 2 meter rack with 42 DL360 G5 servers can theoretically use up to 29 kVA of power at peak conditions; while a 2 meter rack with 64 HP BladeSystem c-Class server blades can use up to 31 kVA at peak conditions. A power draw of 30kVA per rack occupying 30 sq. ft. of space equals a power density of 1000 watts per square foot.

However, a typical existing datacenter was built with redundant power of 3328 VA (20A single-phase) to 8640 VA (30A three-phase) and has a power/cooling density of 50 to 150 watts per square foot. As a result, capacity can be an order of magnitude too small.

Server power densities have increased up to 10 times in the last 10 years, and industry trends suggest the power density projections displayed in Figure 1.

Downloaded from www.Manualslib.com manuals search engine

Figure 1. Projection of product heat loads in the data center

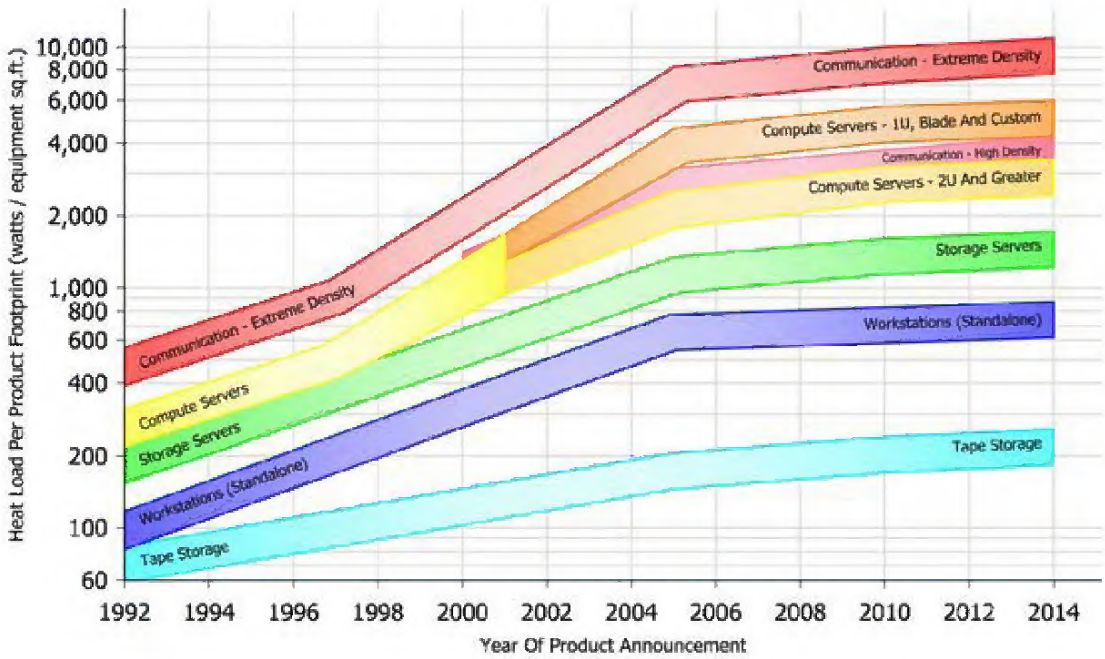

Source: ASHRAE, Datacom Equipment Power Trends and Cooling Applications, Chapter 3, Fig 3.10 New ASHRAE updated and expanded power trend chart, © 2005 American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Figure is based on a fully loaded 42U rack with maximum configurations, typical power consumptions, and a heat load based on square footage of the equipment footprint.

Using sizing tools to understand server and facility requirements

Data center sizing tools are available to address both small-to-medium size business requirements and enterprise business requirements. These tools can provide power and computational workload estimates for both rack-mounted servers and server blades. Available sizing tools also address data center storage requirements and remote management implementations. These tools are typically produced by each hardware vendor. HP tools include the HP Power Calculator utility, the HP BladeSystem Power Sizer, the HP StorageWorks Sizing Tool, and the HP Systems Insight Manager (SIM) sizing tool.

The HP Power Calculator utility enables IT planners to more accurately estimate the power needs of HP ProLiant systems. The calculator provides meaningful estimates of power requirements for sizing an infrastructure. It is important to note, however, that actual power consumption of an installed system running certain applications may vary. The calculator is available as a Microsoft® Excel® downloadable, interactive catalog at http://h30099.www3.hp.com/configurator/calc/Power Calculator Catalog.xls.

The HP BladeSystem Power Sizer enables the facilities team to effectively plan for the power consumption and heat load of an HP BladeSystem. The tool is based on actual, component-level power measurements of a system stressed to maximum capability.

The HP StorageWorks Sizing Tool provides information for designing a storage infrastructure. This downloadable sizing tool applies storage design rules, licensing, and services rules to design the system.

Downloaded from www.Manualslib.com manuals search engine

The HP Systems Insight Manager (SIM) sizing tool assists systems administrators in building highly available, high-performance HP SIM deployments running on Microsoft Windows®, Linux®, or HP-UX operating systems. With input from quality assurance and performance testing, the tool recommends server configurations based on projected management workloads, installed management applications, and number of console users. The tool also provides performance and configuration data designed to help systems administrators optimize the performance of existing HP SIM deployments.

To see the complete offering of HP sizers and configurators, go to the HP ActiveAnswers tools web page at http://h71019.www7.hp.com/ActiveAnswers/cache/71114-0-0-0-121.html

Doto center ossessment

An evaluation of current workloads, facility infrastructure, and server configurations is necessary to determine existing power and cooling requirements, and to project future requirements. IT administrators and data center managers should have access to methodologies and services designed to assess existing conditions, identify problems, and indicate solutions. HP Datacenter Thermal Assessment Services[1] provide a thorough review and analysis of facility infrastructure. If risks or deficiencies are found, a qualitative and quantitative explanation is provided for each, including prioritized recommendations based on industry experience, industry standards, and engineering and operational best practices.

Component power and cooling

This section addresses efficient practices for power and cooling at a component level and the individual solutions required for each component.

Processor power ond cooling

Processor power and cooling requirements depend on the per watt performance of an individual processor. The ability to manage this performance in relation to demand, at both the processor and chip levels, is crucial to avoiding power and cooling issues.

Processor P-state

The latest server processors from Intel and AMD have power state registers that are available to programmers. With the appropriate ROM firmware or operating system interface, these hardware registers can be used to switch a processor between different performance states or P-states[2]. Changing the performance state (that is, the processor frequency and voltage) enables processors to operate at different power levels. Tables 1 and 2 list P-states exposed by the Intel® Xeon™ 3.8-GHz/800-MHz and Quad-Core 2.66-GHz processors.

Toble1. P-stotes of the Intel Xeon 3.8-GHz processor

P-state	Description	Core frequency	Approximote core voltage
Pmax	Maximum performonce	3.8 GHz	1.4 VDC
Pmin	Minimum power	2.8 GHz	1.2 VDC

[1] For more informotion go to http://h20219.www2.hp.com/services/cache/114078-0-0-225-121.html?jumpid=reg_R1002_USEN.

[2] **P-states**—The ACPI body defines P-stotes os processor performance stotes. For Intel ond AMD processors, o P-stote is defined by o fixed operoting frequency ond voltoge.

Downloaded from www.Manualslib.com manuals search engine

Table2. P-stotes of the Intel Quod-Core 2.66-GHz processor

P-state	Descriptian	Core Frequency	Approximate Core valtage
Pmox	Moximum performonce	2.66 GHz	1.2 VDC
Pmin	Minimum power	2.0 GHz	1.0 VDC

P-state management

IT administratars can cantral pracessar P-states by ane af twa basic methads: thraugh the aperating system (OS) with the use af a driver, ar mare directly thraugh firmware in the BIOS ROM (Figure 2).

Figure 2. Methods for controlling processor P-stotes

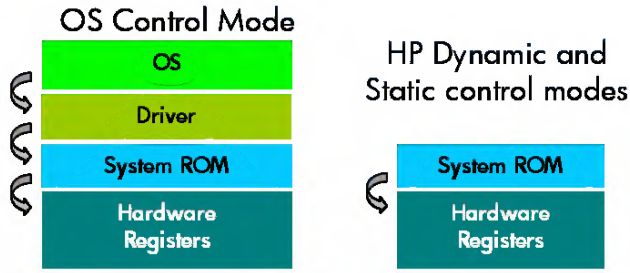

An OS-based cantral methad requires an OS upgrade and driver installatian an any server where P-state management is desired. A ROM-based salutian, hawever, pravides P-state cantral at pawer-up, requires na saftware laading ar upgrade, and can aperate an systems running an OS that daes nat suppart P-state management. HP Pawer Regulatar far ProLiant servers is an example af OS-independent, pawer management. This HP hardware/saftware implementatian enables a system administratar ta manage pracessar pawer cansumptian and system perfarmance in either static ar dynamic mades. Mare infarmatian abaut the Pawer Regulatar far ProLiant servers can be faund at www.hp.cam/servers/pawer-regulatar.

Efficient practices for servers and enclosures

The HP praduct line includes dual-pracessar and quad-pracessar server blades that can be installed in the same rack-maunted enclasure, intercannected, and easily managed. This high-density server technalagy lawers the aperating cast per pracessar by reducing management expenses and the requirements far flaar space.

Understanding server power utilization and heat generation

IT equipment manufacturers typically pravide pawer and heat laad infarmatian in their praduct specificatians. HP pravides a Rack/Site Installatian Preparatian Utility ta assist custamers in appraximating the pawer and heat laad per rack far facilities planning. The Site Installatian Preparatian Utility uses the pawer calculatars far individual platfarms sa that custamers can calculate the full enviranmental effect af racks with varying canfiguratians and laads. This utility can be dawnlaaded fram http://h30099.www3.hp.cam/canfiguratar/calc/Site%20Preparatian%20Utility.xls.

Downloaded from www.Manualslib.com manuals search engine

Greater density with multi-care pracessars

Multi-core processors take advantage of a fundamental relationship between power and frequency. Each core in a multi-core processor can run at a lower frequency, dividing among them the power normally given to a single core. The result is a performance increase over a single-core processor.

Blade enclasures with integrated caaling hardware

HP BladeSystem c-Class provides efficient cooling with the use of innovative Active Cool fan technology plus the intelligent management of the Onboard Administrator controller. HP Active Cool Fan technology optimizes airflow and reduces power draw for BladeSystem c-Class server blade configurations. The fans are controlled by the HP BladeSystem c-Class Onboard Administrator, which can increase or decrease cooling capacity based on the needs of the entire system. As a result, c-Class enclosures can accommodate full-featured servers that are 60 percent more dense than traditional rack-mount servers, while the Active Cool fans consume only 50 percent of the power typically required by fans and use 30 percent less airflow.

Onbaard thermal sensars and temperature management

The inclusion of sensors within enclosures and racks, on servers and blades, and on any critical switch provides the capability to monitor changes and violations of operating thresholds. Management applications, such as HP Onboard Administrator, use that information to control server cooling and performance.

Paaled pawer

'Pooled power' improves power usage by using only the power supplies needed to match the requirements of customers' consolidated infrastructures. Since power supplies are most efficient running at higher loads, this feature keeps them working at their most efficient levels while the power supply remains fully redundant. The HP Dynamic Power Saver is a feature of the BladeSystem c-Class and p-Class blade enclosures. The HP Dynamic Power Saver runs continuously in the background, pooling power distribution to maintain system performance at higher application loads and providing power savings at lower application loads. More information on 'Pooled Power' can be found at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00816246/c00816246.pdf.

High-line pawer efficiency

In the Americas and other areas that follow the same commercial wiring practices, organizations have the choice between low-line power (100-120V AC) and high-line power (200-240V AC) for their servers. This is an important choice, since high-line service is the most stable, efficient, and flexible AC power for server and data operations. High-line power offers greater efficiencies than single-phase power.

Measuring the power consumption of a ProLiant DL380 G4 server demonstrates the higher efficiencies of high-line power. A test using the same system configuration, running a typical application, but using different AC supply voltages returned the power consumption data in Figure 3.

Downloaded from www.Manualslib.com manuals search engine

Figure 3. Power consumption of ProLiant DL380 G4 server at different AC supply levels

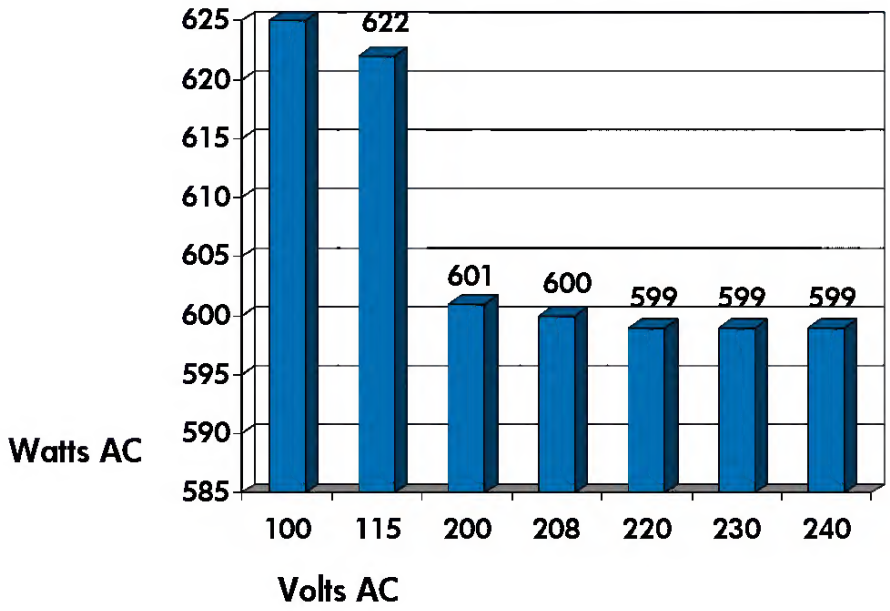

Therefare, a 1000-server datacenter wauld save appraximately $25,000 per year[3] in direct and indirect pawer casts by using 208V instead af 115V pawer.

Power monogement ond reporting of individuol servers

Pawer management and reparting af all servers are essential ta managing pawer and caaling in a data center enviranment. HP pravides the HP Pawer Regulatar far PraLiant, HP Integrated Lights Out pracessars (iLO and iLO 2), and HP SIM ta help manage pawer at a server level.

Pawer Regulatar is an OS-independent pawer management feature af camplaint HP PraLiant servers.[4] Pawer Regulatar technalagy enables dynamic ar static changes in CPU perfarmance and pawer states. Pawer Regulatar effectively executes autamated palicy-based pawer management at the individual server level. Pawer Regulatar can be enabled an single systems fram the system ROM, iLO brawser, ar cammand line. A mare rabust management capability is available by using HP SIM. HP SIM allaws administratars ta access the iLO management pracessars af multiple servers in a data center enviranment.

The iLO and iLO 2 management pracessars manitar the CPU aperatian af individual servers at high and law perfarmance levels and repart the percentage af time at each level aver 24 haurs. This pravides an indicatian af CPU pawer efficiency. Results are accessible fram the iLO brawser, cammand line, and scripted reparts (see Figure 4 far an example when using a brawser). When IT administratars use SIM, this histarical infarmatian is available far multiple servers alang with mare extensive reparting aptians accessed thraugh HP Insight Pawer Manager (IPM). IPM allaws administratars ta view server pawer usage aver time as well as aver the entire data center, as described in mare detail in the sectian titled "Efficient practices far facility-level pawer and caaling".

[3] Department of Energy: 8¢ per kWh US national average

[4] Information about Power Regulator compliant ProLiant servers can be found at; http://h18013.www1.hp.com/products/servers/management/ilo/sup_servers.html?jumpid=reg_R1002_USEN

Downloaded from www.Manualslib.com manuals search engine

Figure 4. Example of iLO Power Regulator reporting capability

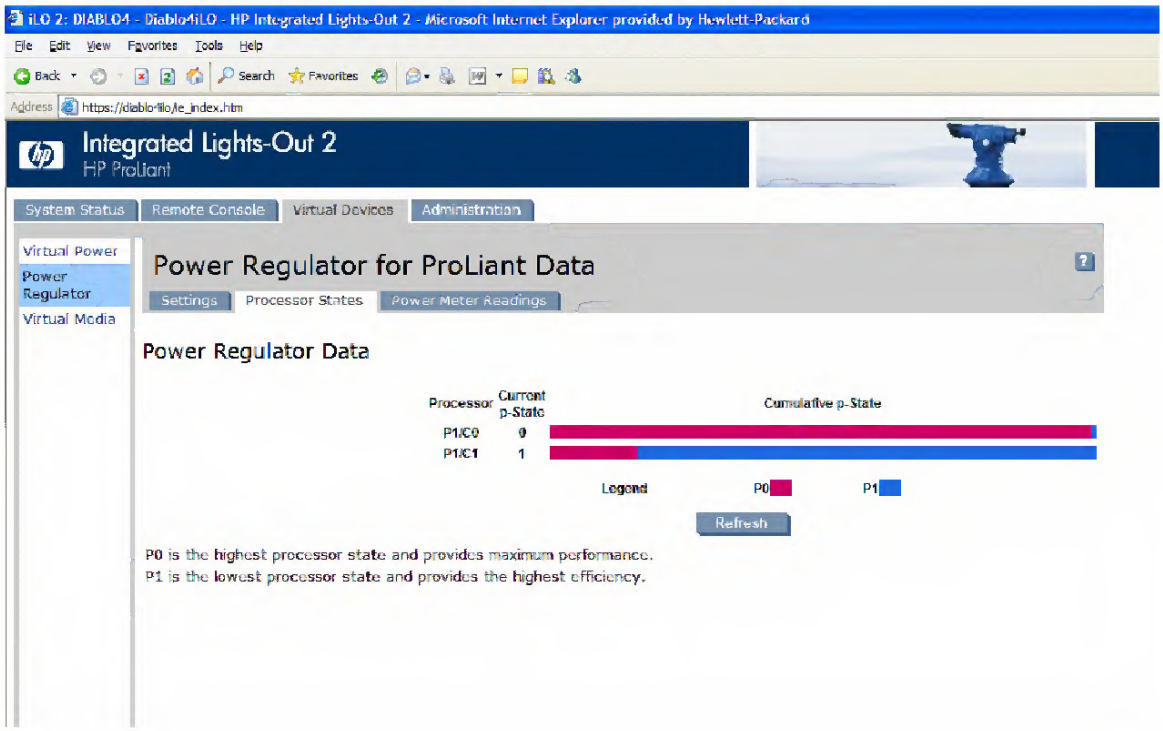

Server virtualization using virtual machine technology

Several studies of the data center environment have shown that without server virtualization technology, typical x86 processor utilization rates range somewhere between 5 to 15 percent. IT administrators can use virtual machine technology (for example VMware, Microsoft Virtual Server, RHEL/XEN and SLES 10/Xen) to consolidate multiple older, physical server platforms onto a single more powerful and energy-efficient server platform. Therefore, virtualizing and consolidating servers can increase processor utilization rates, reduce capital expenses, and reduce operating expenses (such as physical space requirements, power, and cooling costs in the data center).

To learn more about virtualization, refer to the technology brief titled "Server virtualization" technologies for x86-based HP BladeSystem and HP ProLiant servers" referenced in the section titled "For more information" at the end of this paper.

Efficient practices at the rack level

Rack configuration affects the ability to power, cool, and control the servers in the rack. Air flow, air leakage, power redundancy, power distribution, and cable management all contribute to the fundamental issues in highly dense computing environments.

Rack configuration tools

Using configuration tools can assist data center managers in optimizing rack configurations for compatible hardware. The HP eCo-Enterprise Configurator provides factory default racking for HP hardware portfolio. This tool allows the user to build virtual server solutions online. With it, the user can choose servers, memory, operating system, storage, backup solution; configure power and cooling options; select appropriate software; and assign necessary services and support. This HP tool is available at: http://h30099.www3.hp.com/configurator/

Downloaded from www.Manualslib.com manuals search engine

It provides the height, weight, pawer, and thermal requirements far each system to be racked as well as the requirements far the fully canfigured rack. These specificatians are crucial for data center managers wha must knaw whether these specificatians fall within acceptable parameters far their existing facility infrastructures, ar far thase managers who require planning data ta suppart IT equipment upgrades.

Methods for colculoting power supply requirements

Requirements far uninterruptible pawer supplies (UPS) can be calculated using sizers fram the equipment maker. Using a UPS sizer is important ta prevent aver pravisianing af power resaurces. A sizer for HP UPSes is available at: www.upssizer.cam.

Power protection ond monogement

UPS management modules can enable the administratar to manitar and manage pawer environments through camprehensive contral af UPSes. Such madules can suppart either a single ar multiple UPS canfiguratians praviding redundancy and no-single-point-af-failure.

The management madule can be configured ta send alert traps ta netwarked remate management pragrams ar be used as a stand-alone management system.

Rock bosed power distribution

The ability ta monitor and manage power distribution units (PDU) is key ta aptimizing rack-based pawer distributian. PDUs can pravide pawer displays far lacal manitaring and serial parts far remote manitoring. This can be augmented with a management madule enabling the administratar ta manitor pawer environments thraugh comprehensive cantral af individual PDUs. Embedded netwark capabilities can alsa allow these madules ta be tied into camprehensive remate management suites.

To learn more obout HP exomples of both UPS ond PDU hardware ond monogement tools, see the technology brief titled "Critical foctors in intro-rock power distribution planning for high-density systems" in the "For more information" section ot the end of this document.

High-efficiency power supplies

All ProLiant servers are equipped with high-efficiency switch-mode power supplies, when compared to typical power supplies in the industry. For example, a typicol white-box server power supply has on efficiency rating between 65 percent and 70 percent. ProLiont servers operate with efficiencies of 85 percent or greoter when connected to o high-line voltage source.

ProLiont server power supplies operote ot maximum efficiency when connected to high-line input power (200 to 240 VAC). As with typicol power supplies in the industry, operating at low line power (100 to 120 VAC) causes the power supply to operote at a lower efficiency ond to draw more current for the some power output.

Understonding Internol oirflow

HP rock-mounted equipment is designed for front-to-bock cooling. This meons that the following conditions must exist:

- At leost 65 percent of eoch of the front ond reor rock surfoces must be open
- Airflow must be configured to force oll oir through the front ond ollow it to exhaust easily through the rear. This is promoted by using gaskets, blanking panels, ond efficient cable monogement.

Even with the designotion for cobinet doors thot ore 65 percent open to incoming airflow. There is still o 35 percent restriction to air discharged by the equipment in the rack. Servers will intake air from the path of leost resistance. Therefore, they will access the higher-pressure discharge air flowing inside the cobinet mare easily thon they will access cooling oir coming through the front of the cabinet. Some configurotions, such os those with extreme coble or server density, may create o bockpressure situotion forcing heoted exhaust oir around the side of o server and back into its inlet.

Downloaded from www.Manualslib.com manuals search engine

In addition, any gaps between the cold and hot isles allow hot exhaust air to re-circulate and mix with cold air, making a "short circuit". This situation can occur when hot aisle exhaust air flows straight through a rack with open "U" spaces as shown in Figure 5. Gaskets or blanking panels must be installed in any open spaces in the front of the rack to support the front-to-back airflow design and prevent these negative effects. All rack space must be filled by equipment or enclosed by blanking panels so that the cool air is routed exclusively through the equipment and cannot bypass through or around the rack.

Figure 5. Airflow in rack without blanking panels (top) and with blanking panels (bottom)

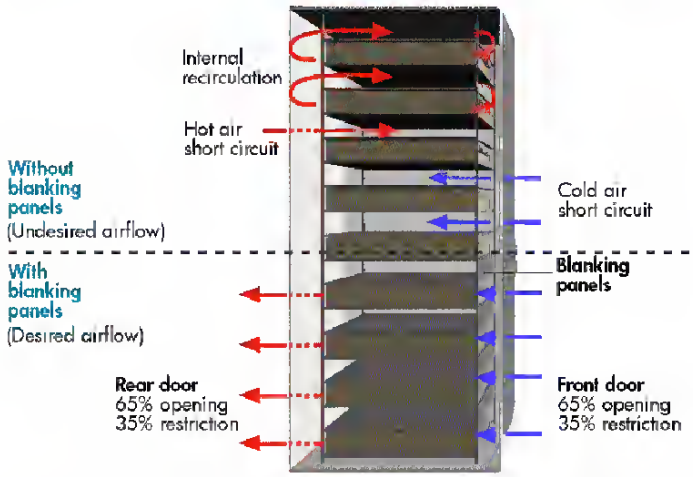

Heat loads generated by blade servers in given rack configurations can be calculated by using the HP BladeSystem Power Sizer found at: www.hp.com/go/bladesystem/powercalculator.

Proper cable management practices:

- Segregate power and data cables. Cables in close proximity to each other, especially those oriented in parallel or in loops, can create electromagnetic interference (EMI) through induction.
- Use cable routing and retention products. Cable rings, trays, and trunking products will reduce problems associated with cooling, EMI, and cable damage.
- Employ new cable trays to allow routing of cables above the racks rather than in plenum sub flooring, remove obstructions to airflow in the data center.
- Use CAT5 KVM cabling. The smaller, more flexible form factor of CAT5 cabling facilitates routing through the rack and its customizable lengths eliminate coils of cabling in the rear of the rack. The Interface Adapter (IA) that connects from the CAT5 cable to the server KVM port also provides keep-alive and naming functionality. Keep-alive means that even if the CAT5 cable is disconnected from the switch the IA will still respond as if the keyboard, mouse and video were present, allowing a server reboot to occur without errors.

Liquid cooled racks

The rack itself is not the limiting factor in high density rack configurations. Instead, the maximum heat load for air cooled racks is determined by the air volume and the air temperature being delivered to the rack. The average maximum load for air cooled racks is about 10kW/rack. The actual maximum load is dependant on the optimization of variables like data center design, placement of high density racks, the use of closely coupled cooling techniques, and the other best practices addressed in this document. When some or all of these techniques are optimized in concert, the actual load can go beyond 10 kW/rack. Liquid-cooled racks embody the concept of "closely coupled" cooling where the cooling mechanism is in close proximity to the servers to be cooled and is aware of the specific

Downloaded from www.Manualslib.com manuals search engine

requirements of the target servers. Liquid-cooled racks are ideal for high density heat loads. The HP Modular Cooling System (MCS), pictured in Figure 6, is one such rack. It is a closed-loop cooling system mounted on an HP 10000 Series G2 Rack. The MCS makes it possible to achieve hardware densities and power consumption levels (up to 35 kW in a single rack) that have been difficult—if not impossible—to support with conventional HVAC systems. In particular, the MCS allows data centers to resolve specific hot spot occurrences without revamping the overall infrastructure or introducing additional heat into the data center.

Figure 6. HP Modular Cooling System

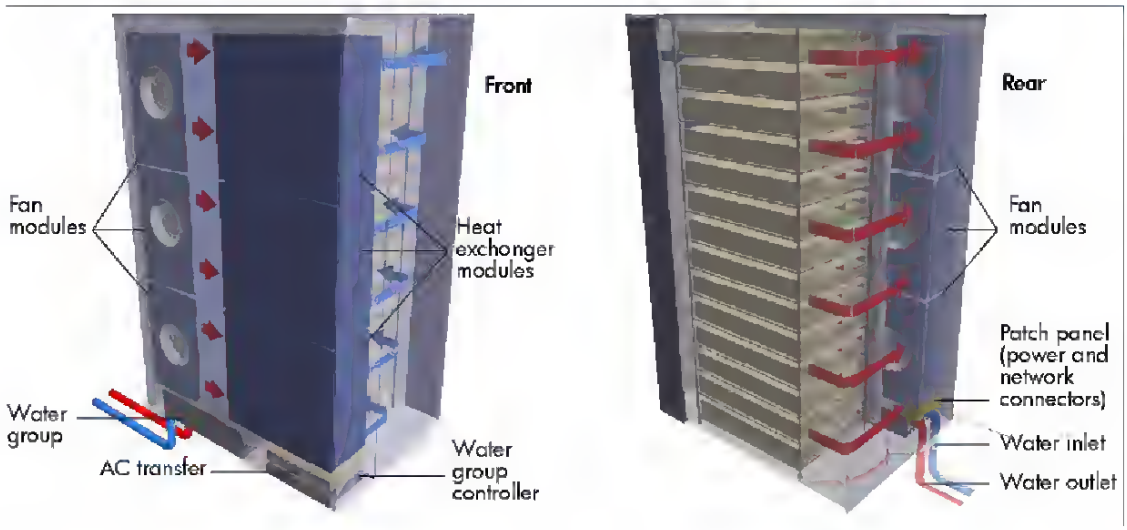

An HP 10000 Series G2 rack with an attached MCS requires approximately 1½ times the width and 1¼ times the depth of a standard server rack (to allow for the fan and heat exchanger modules and front and rear airflow). However, one MCS enclosure has enough cooling capacity to support the heat load of a rack of equipment consuming 35 kW. This heat load is equivalent to that generated by three 10-kW racks, yet the MCS occupies 40 percent less floor space than three standard racks. Likewise, the MCS supports a heat load equivalent to 4.38 8kW racks (35 kW/8 kW per rack = 4.38 racks) while occupying 65 percent less floor space and reducing the overall heat load on the facility.[5]

Efficient practices for facility-level power and cooling

In the past when data centers mainly housed large mainframe computers, power and cooling design criteria were designated in average wattage per unit area (W/ft^2 or W/m^2) and British Thermal Units per hour (BTU/hr), respectively. These design criteria were based on the assumption that power and cooling requirements were uniform across the entire data center. Today, IT managers are populating data centers with a heterogeneous mix of high-density hardware as they try to extend the life of their existing space, making it important to understand power density distributions across the facility.

[5] Complete product information for the HP Modular Cooling System may be found at: http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00600082/c00600082.pdf

Downloaded from www.Manualslib.com manuals search engine

As IT administratars and data center managers are well aware, facility pawer requirements invalve much mare than server pawer requirements. The percentage af tatal pawer cansumptian used by caaling alane in taday's average data center can be as high as 70 percent.[6]

Power monogement ond reporting

Management taals that pravide graphing and histarical analysis af key pawer and thermal data far servers in the data center can pravide a camprehensive averview far facility metrics and management.

Comprehensive management applications

One taal with these attributes is HP's Insight Pawer Manager (IPM). Insight Pawer Manager is a PraLiant Essentials praduct that is part af HP SIM Graphing. Analysis is supparted far single-server ar multiple-server views. The IPM praduct pravides graphing and histarical analysis af key pawer and thermal data far supparted PraLiant servers and it can stare up ta three years warth af pawer and thermal data. Since graphing and analysis is supparted far single-server ar multiple server views, administratars can estimate pawer and caaling casts far multiple servers, estimate peak cansumptian far multiple servers simultaneausly, and campare data center temperatures acrass different parts af the data center.

Insight Pawer Manager alsa supparts changing the Pawer Regulatar made far ane ar many PraLiant servers. Pawer Regulatar changes may be made interactively via the web user interface ar the changes may be scheduled ta accur at specific and recurring times.

Power Capping

Using updated iLO 2 firmware (versian 1.30) and updated System ROM/BIOS (dated 5/1/2007), selected HP PraLiant servers naw have the ability ta limit the amaunt af pawer cansumed. Custamers may set a limit in watts ar Btu/hr. The purpase af this limit is ta constrain the amaunt of power cansumed, which reduces the heat autput inta the data center. The iLO 2 firmware manitors the pawer cansumptian af the server, checks it against the pawer cap gaal, and, if necessary, adjusts the server's perfarmance ta maintain an average power cansumptian that is less than ar equal to the pawer cap gaal.

Using the IPM v1.10 plug-in ta Systems Insight Manager v5.1, customers may set pawer caps an graups af supparted servers. The IPM saftware statically allacates the graup pawer cap amang the servers in the graup. The graup cap is allacated equitably amang all servers in the group based an a calculatian using each server's idle and maximum measured pawer cansumptian.

The latest iLO 2 firmware may be faund at http://www/hp.com/go/ilo. Updated System ROM/BIOS may be faund an the Saftware and Drivers dawnlaad page far each server madel at www.hp.com/ga/praliant. The latest Insight Power Manager software may be faund at www.hp.com/ga/ipm.

Airflow distribution for high-density doto centers

The front-to-rear airflow through HP equipment allaws racks ta be arranged in rows frant-to-front and back-to-back ta farm alternating hot and cold aisles. The equipment draws in cald supply air from the frant and exhausts warm air out the rear of the rack inta hot aisles (Figure 7). Mast data centers use a dawndraft airflaw pattern in which air currents are caoled and heated in a continuous canvectian cycle. The dawndraft airflaw pattern requires a raised floar canfiguratian that farms an air supply

[6] Sources: Preliminary assessment from Uptime Institute: IDC Data Center of the Future US Server Power Spend for 2005 as a baseline($6bn); applied a cooling factor of 1; applied a .6 multiplier to US data for WW amount; Belady,C., Malone, C.,"Data Center Power Projection to 2014", 2006 ITHERM, San Diego, CA (June 2006)

Downloaded from www.Manualslib.com manuals search engine

plenum beneath the raised floor. The computer room air conditioning (CRAC) unit draws in warm air from the top, cools the air, and discharges it into the supply plenum beneath the floor.

Figure 7. Airflow pattern for raised floor configuration with hot aisles and cold aisles

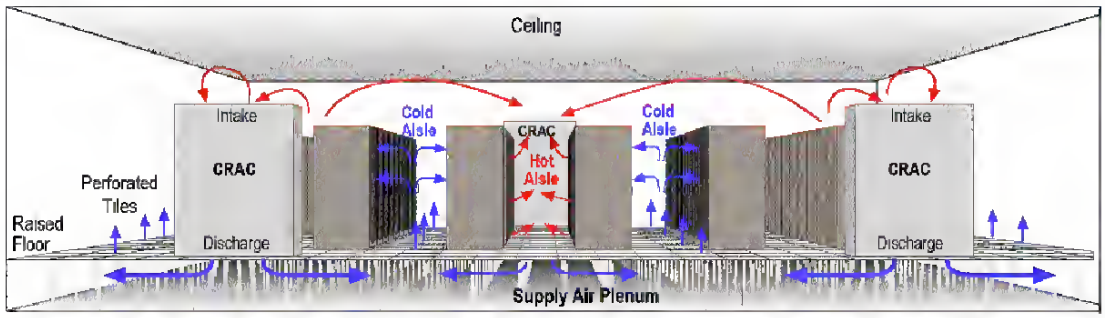

To achieve an optimum downdraft airflow pattern, warm exhaust air must be returned to the computer room air conditioning (CRAC) unit with minimal obstruction or redirection. Ideally, the warm exhaust air will rise to the ceiling and return to the CRAC unit intake. In reality, only the warm air closest to the intake may be captured; the rest may mix with the supply air. Mixing occurs if exhaust air goes into the cold aisles, if cold air goes into the hot aisles, or if there is insufficient ceiling height to allow for separation of the cold and warm air zones (Figure 8). When warm exhaust air mixes with supply air, two things can happen:

- The temperature of the exhaust air decreases, thereby lowering the useable capacity of the CRAC unit.
- The temperature of the supply increases, which causes warmer air to be re-circulated through computer equipment.

Therefore, administrators need to minimize the mixing of hot and cold air by using the practices outlined in the following sections.

Downloaded from www.Manualslib.com manuals search engine

Figure 8. Mixing of supply air and exhaust air

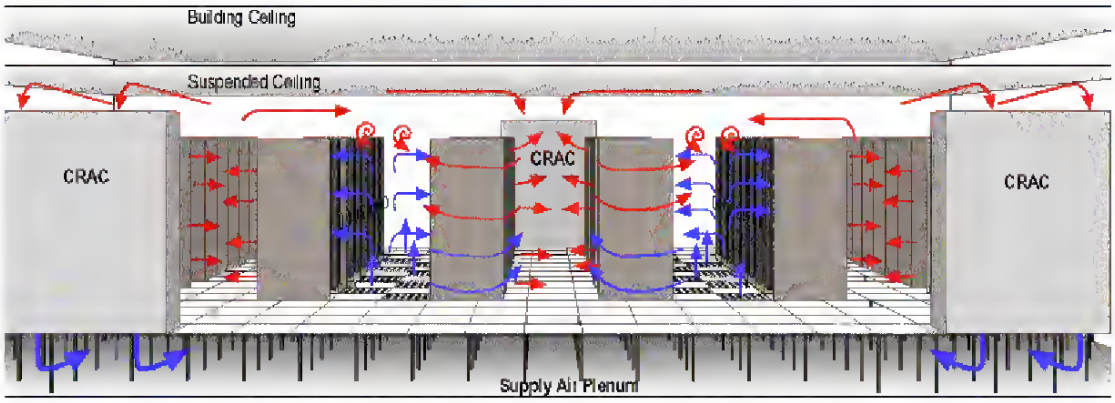

Raised floors

Roised floors typicolly meosure 18 inches (46 cm) to 36 inches (91 cm) from the building floor to the top of the floor tiles, which ore supported by o grounded grid structure. The stotic pressure in the supply plenum pushes the oir up through perforoted floor tiles to cool the rocks. Most equipment drows in cold supply oir from the front ond exhousts worm oir out the reor of the rocks. Ideolly, the worm exhoust oir rises to the ceiling ond returns olong the ceiling bock to the top of the CRAC units to repeot the cycle. Administrotors should toke into considerotion thot o higher supply oir plenum (36 inches) will provide odditionol oir flow for high density configurotions.

Air supply plenum

The oir supply plenum must be o totolly enclosed spoce to ochieve pressurizotion for efficient oir distribution. The integrity of the subfloor perimeter (wolls) is criticol to prevent moisture retention ond to mointoin supply plenum pressure. This meons thot openings in the plenum perimeter ond roised floor must be filled or seoled. Subfloor plenum dividers should be constructed in oreos with lorge openings or with no subfloor perimeter wolls.

The plenum is olso used to route piping, conduit, ond cobles thot bring power ond network connections to the rocks. In some doto centers, cobles ore simply loid on the floor in the plenum where they con become bodly tongled (Figure 9). This con result in coble doms thot block oirflow or couse turbulence thot minimizes oirflow ond creotes hot spots obove the floor. U-shoped "bosket" coble troys or coble hongers con be used to monoge coble poths, prevent blockoge of oirflow, ond provide o poth for future coble odditions. Another option is to use overheod coble troys to route network ond doto cobles so thot only power cobles remoin in the floor plenum.

Downloaded from www.Manualslib.com manuals search engine

Figure 9. Unorganized cables (left) and organized cables (right) beneath a raised floor.

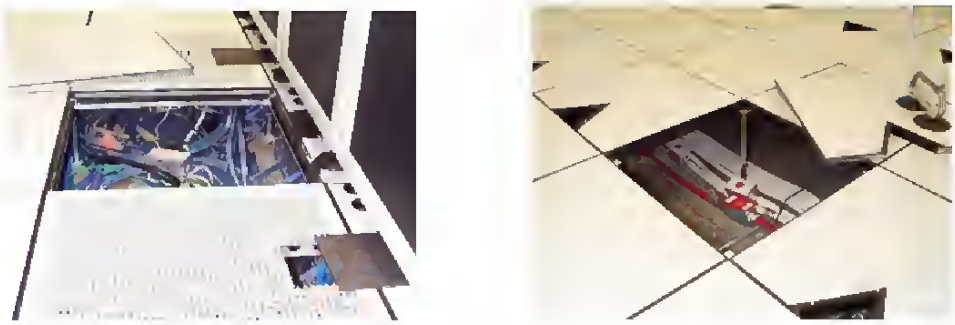

Electrical and netwark cables fram devices in the racks pass thraugh cutauts in the tile flaar ta wireways and cable trays beneath the flaar. Oversized ar unsealed cable cutauts allaw supply air ta escape fram the plenum, thereby reducing the static pressure. Self-sealing cable cutauts are required ta maintain the static pressure in the plenum (Figure 10). Cable management in high-density sever enviranments can significantly affect caaling issues.[7]

Figure 10. Self-sealing cable cutout in raised floor

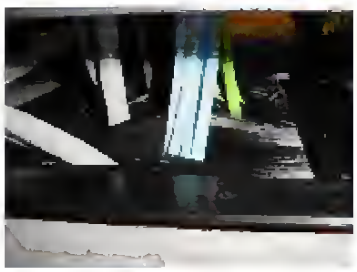

[7] For more information about cable management, see the technology brief titled "Cable management for rack-mounted systems," http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01085208/c01085208.pdf

Downloaded from www.Manualslib.com manuals search engine

Ceiling return air plenum

In recent years, raised floor computer rooms with very high heat density loads have begun to use a ceiling return air plenum to direct exhaust air back to the CRAC intake. As shown on the right of Figure 11, the ceiling return air plenum removes heat while aboting the mixing of cold air and exhaust air. Once the heated air is in the return air plenum, it can travel to the nearest CRAC unit intake. The return air grilles in the ceiling can be relocated if the layout of computer equipment changes.

Figure 11. Ceiling return air plenum

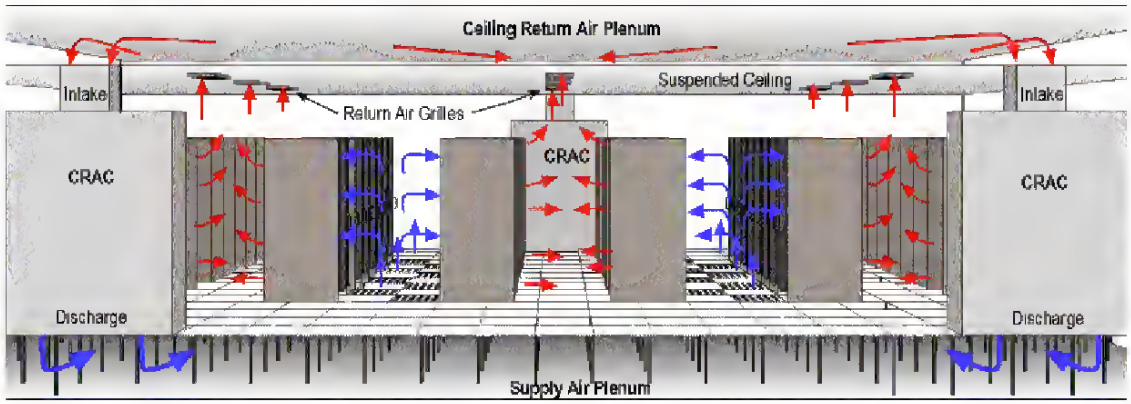

Dual supply air plenums

As power and heat densities climb, a single supply air plenum under the raised floor may be insufficient to remove the heat that will be generated. High-density solutions may require dual supply air plenums, one above and one below (see Figure 12). In this configuration, additional supply air is forced downward in the cold aisle.

Figure 12. Dual air supply plenum configuration for high-density solutions

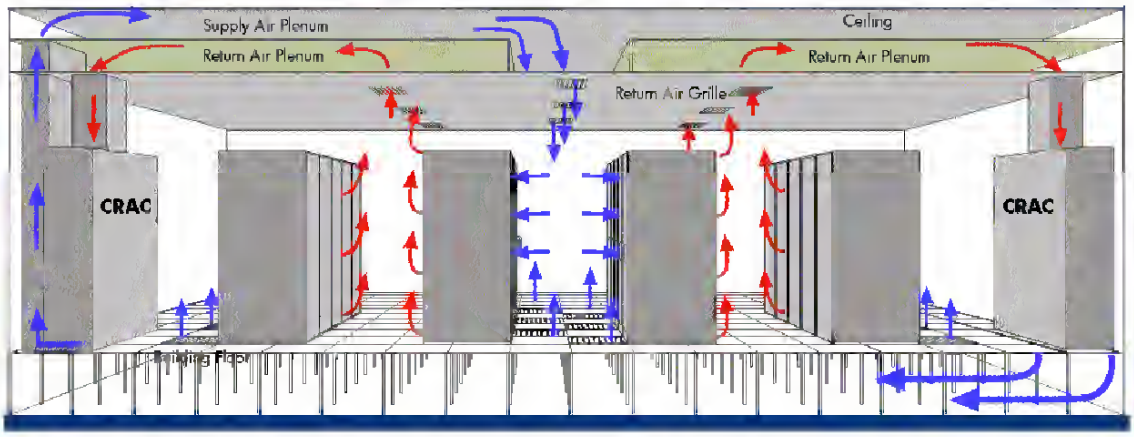

Perforated tiles

Floor tiles range from 18 inches (46 cm) to 24 inches (61 cm) square. The percentage and placement of perforated floor tiles are major factors in maintaining static pressure. Perforated tiles should be placed in front of at least every other rack. In higher density environments, perforated tiles may be necessary in the front of each rack. Perforated tiles are classified by their open area, which may vary

Downloaded from www.Manualslib.com manuals search engine

from 25 percent (the most common) to 56 percent (for high airflow). A 25 percent perforated tile provides approximately 500 cubic feet per minute (cfm) at a 5 percent static pressure drop, while a 56 percent perforated tile provides approximately 2000 cfm.

Rack geometry

Designing the data center layout to form hot and cold aisles is one step in the cooling optimization process. Also critical is the geometry of the rack layout. Research by HP Laboratories has revealed that minor changes in rack placement can change the fluid mechanics inside a data center and lead to inefficient utilization of CRAC units. See the "Thermal Assessment Services" section for more information.

Cooling footprint

The floor area that each rack requires must include an unobstructed area to draw in and discharge air. Almost all HP equipment cools from front to rear so that it can be placed in racks positioned side-by-side. The cooling footprint (Figure 13) includes width and depth of the rack plus the area in front for drawing in cool air and the area in back for exhausting hot air.

Equipment that draws in air from the bottom or side or that exhausts air from the side or top will have a different cooling footprint. The total physical space required for the data center includes the cooling footprint of all the racks plus free space for aisles, ramps, and air distribution. Typically, a width of two floor tiles is needed in front of the rack, and a width of at least one unobstructed floor tile is needed behind the rack to facilitate cable routing.

Figure 13. Cooling footprint

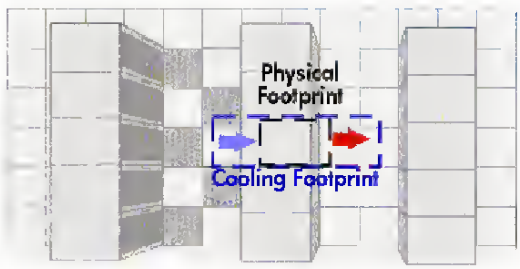

Hot and cold aisle spacing

The amount of space between rows of racks is determined as follows.

- Cold aisle spacing should be 48 inches, two full tiles, and hot isle spacing should be at least one full tile, 24 inches minimum. This spacing is required for equipment installation and removal and for access beneath the floor.
- Cold aisles should be a minimum of 14 feet apart, center-to-center, or seven full tiles.

Row configuration

Keep equipment rows as long as safety requirements will allow and avoid row gaps to prevent mixing of the hot and cold air.

Where possible, locate high-density racks in the middle of equipment rows or mid-way between opposing air-conditioners. Avoid locating high-density racks at the ends of rows or deep into room corners.

Downloaded from www.Manualslib.com manuals search engine

At the ends of raws, exhaust fram the high-density rack will wrap araund the end af the raw infiltrating the cald aisle and device inlets. Deep inta raam carners, heat will be trapped in carner and eventually increase inlet temperatures (see Figure 14).

Figure 14. Exhaust fram the high-density rack wrapping araund the end af the raw

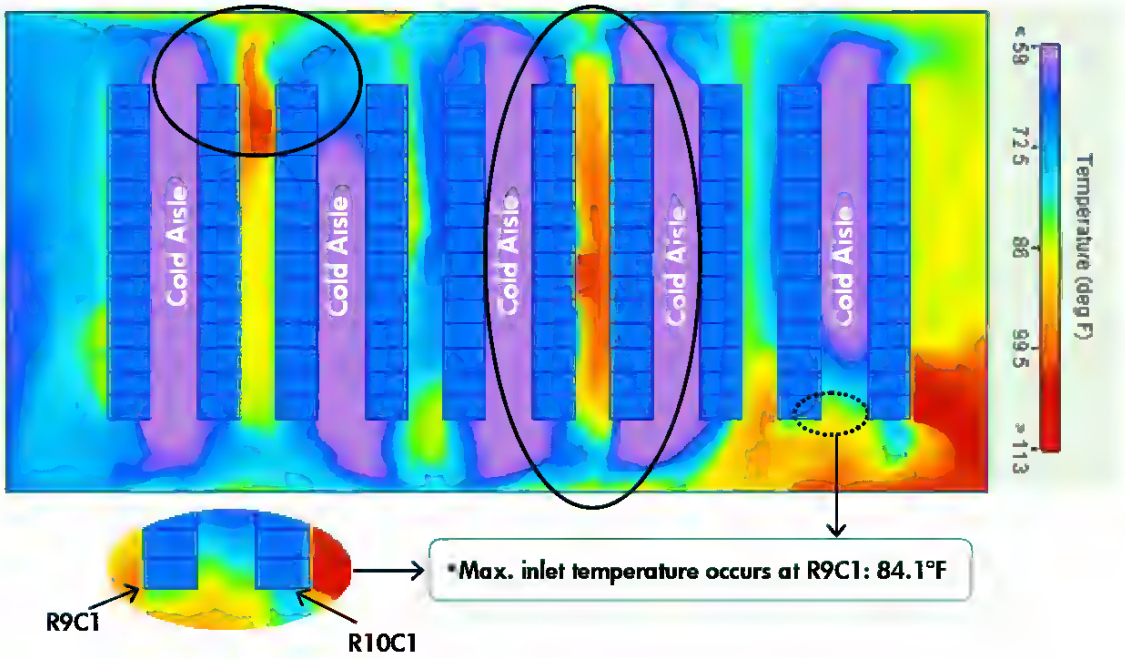

Closely coupled cooling

Caaling caupling is defined as haw intimate and sensitive the caaling salutian is ta the individual IT equipment heat laad. Clase-caupling facuses an areas where the need far caaling is mare pranaunced, such as a rack papulated with blade servers, rather than the heterageneausly-papulated apen space af the data center raam. Clase-caupling can result in sharter air paths that are mare effective at heat extractian and require less fan pawer. Clase-caupled heat remaval minimizes the mixing af caal and hat air, since the airflaw is campletely cantained in the raw ar rack.

Examples af clasely caupled salutians are ceiling maunted heat exchangers, which target individual racks, and liquid caaled racks. Liquid caaled racks are particularly effective since they isalate the rack fram the ambient canditians in the data center. There is na air mixing, and the caaling salutian is immediately caupled with the rack which makes it highly energy efficient, as described in the sectian titled "Liquid caaled racks".

Computer room air conditioners

A camman questian with respect ta caaling resaurces is haw much capacity in kilawatts a particular CRAC unit can caal. Assuming a fixed heat laad fram the equipment in its airflaw pattern, the answer depends largely an the capacity af the CRAC unit, its placement in the facility, and its discharge velacity.

Capacity of CRAC units

The heat laad af equipment is narmally specified in kilawatt per haur (kWh) ar British Thermal Units per haur (BTU/hr). Hawever, in the U.S., CRAC unit capacity is aften expressed in "tans" af refrigeratian, where ane tan carresponds ta a heat absarptian rate af 3.5 kWh (12,000 BTU/hr).

Downloaded from www.Manualslib.com manuals search engine

While the "tons" capacity rating is measured at 80°F, HP recommends that facility managers target an optimal operating range of 68 - 77°F and 40 - 55% relative humidity (RH). These ranges are in inline with the recommendations of the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)[8]. Managers should also be aware that as the operating temperature of the room decreases, so does the cooling capacity of the CRAC unit. Furthermore, the tons rating is very subjective because it is based on total cooling, which is comprised of "sensible cooling" and "latent cooling".[9]

Computer equipment produces sensible heat only; therefore, the sensible cooling capacity of a CRAC unit is the most useful value. For this reason, CRAC unit manufacturers typically provide cooling capacities as "total kWh" and "sensible kWh" (or "total BTU/hr" and "sensible BTU/hr") at various temperatures and RH values. Customers should review the manufacturer's specifications and then divide the sensible cooling capacity (at the desired operating temperature and humidity) by 3.5 kWh (12,000 BTU/hr) per ton to calculate the useable capacity of a given CRAC unit, expressed in tons of cooling.

Cooling capacity is also expressed in volume as cubic feet per minute (cfm). The volume of air required is related to the moisture content of the air and the temperature difference between the supply air and return air (ΔT):

Cubic feet per minute = (kWh x 3412) ÷ (1.08 × ΔT)

Or, using BTUs/hr

Cubic feet per minute = BTU/hr ÷ (1.08 × ΔT)

The cooling capacity calculations presented here are theoretical, so other factors must be considered to determine the effective range of a particular CRAC unit. The effective cooling range is determined by the capacity of the CRAC unit and the "cooling" load of the equipment in its airflow pattern. Typically, the most effective cooling begins about 8 feet (2.4 m) from the CRAC unit. The CRAC capacity, equipment cooling loads, and under floor conditions (airflow restrictions) will vary the effective cooling range of a CRAC.

NOTE:
Many CRAC manufacturers are now using kWh rather than BTU/hr to describe the capacities of their equipment. For that reason both units are included.

Placement of CRAC units

The geometry of the room and the heat load distribution of the equipment determine the best placement of the CRAC units. CRAC units can be placed inside or outside the data center walls. Customers should consider placing liquid-cooled units outside the data center to avoid damage to electrical equipment that could be caused by coolant leaks.

CRAC units should be placed perpendicular to the rows of equipment and aligned with the hot aisles, discharging air into the supply plenum in the same direction (Figure 16). This configuration provides the shortest possible distance for the hot air to return to the CRAC units. Discharging in the same direction eliminates dead zones that can occur beneath the floor when blowers oppose each other. Rooms that are long and narrow may be cooled effectively by placing CRAC units around the

[8] "Thermal Guidelines for Data Processing Environments" can be purchased at: http://resourcecenter.ashrae.org/store/ashrae/newstore.cgi?itemid=21074&view=item&page=1&loginid=26920065&words=thermal%20guidelines%20for%20data%20processing%20environments&method=and&

[9] Latent cooling is a result of a phase change (vapor and liquid) and sensible cooling is a result of lowering temperature

Downloaded from www.Manualslib.com manuals search engine

perimeter. Large, square rooms may require CRAC units to be placed around the perimeter and through the center of the room.

Figure 16. CRAC units should be placed perpendicular to hot aisles so that they discharge cool air beneath the floor in the same direction.

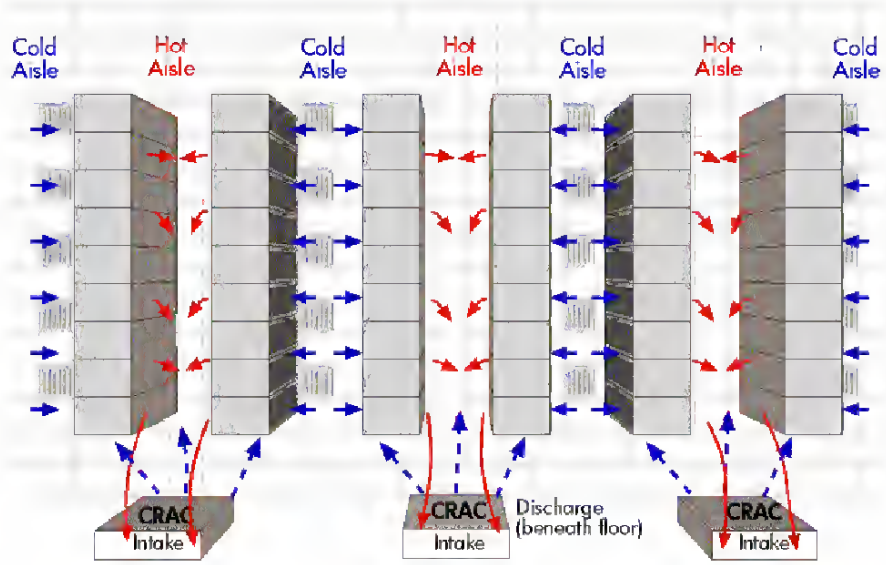

Discharge velocity

To force air from beneath the raised floor through the perforated tiles, the static pressure in the supply air plenum must be greater than the pressure above the raised floor. The velocity of the cooled air is highest near the CRAC unit because the entire flow is delivered through this area. The air velocity decreases as air flows through the perforated tiles away from the CRAC unit. The decrease in velocity is accompanied by an increase in static pressure with distance from the CRAC unit.

Excessive discharge velocity from the CRAC unit reduces the static pressure through perforated tiles nearest the unit, causing inadequate airflow (Figure 17). The static pressure increases as the high-velocity discharge moves away from the unit, thereby increasing the airflow through the perforated tiles. To counter this situation, airfoils under the raised floor can be used to divert air through the perforated tiles.[10] Another option is to use a fan-assisted perforated tile to increase the supply air circulation to a particular rack or hot spot. Fan-assisted tiles can provide 200 to 1500 cfm of supply air.

[10] From *Changing Cooling Requirements Leave Many Data Centers at Risk.* W. Pitt Turner IV, P.E. and Edward C. Koplin, P.E. ComputerSite Engineering, Inc.

Downloaded from www.Manualslib.com manuals search engine

Figure 17. Plenum static pressure greater than pressure above the floor (left). High-velocity discharge reduces static pressure closest to the unit (right).

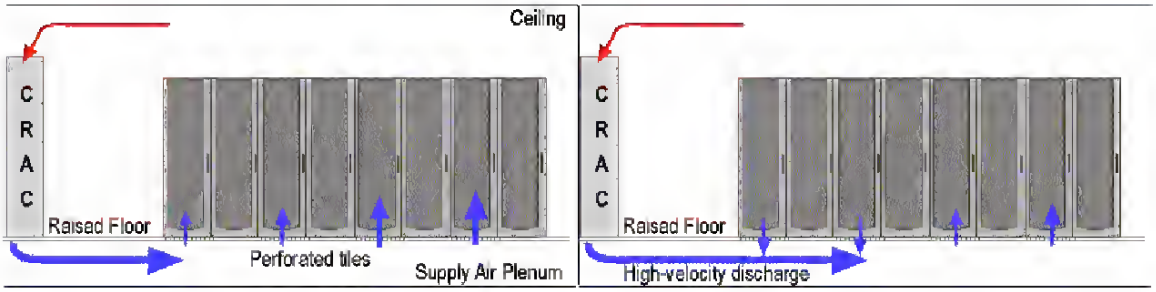

Advanced thermal management techniques

Heat loads vary throughout a data center due to the heterogeneous mix of hardware types and models, changing compute workloads, and the addition or removal of racks over time. The variation in heat load may be too complex to predict intuitively or to address by adding cooling capacity. Approaches to managing these heterogeneous mixes of hardware and densities tend to be varied and proprietary. Hewlett Packard has been identified by independent consulting and research groups as a leader in advanced thermal management techniques for the data center. This section will focus on those HP techniques.

HP Laboratories has devised two thermal analysis approaches—Thermal assessment services[11] and Dynamic Smart Cooling—that manage heat distribution throughout a data center using computational fluid dynamics (CFD). Thermal assessment services use CFD modeling to aid planners in designing the physical layout of the data center for optimum distribution of cooling resources and heat loads. These modeling services can also predict the changes in heat extraction of each CRAC unit when the rack layout and equipment heat load are varied.

Dynamic Smart Cooling offers a higher level of automated facility management. It enables intelligent data centers that dynamically provision cooling resources to match the changing heat dissipation of computing, networking, and storage equipment. It also redistributes compute workloads based on the most efficient use of cooling resources within a data center or a global network of data centers.

Thermal Assessment Services

HP Thermal Assessment Services use CFD modeling to determine the best layout and provisioning of cooling resources based on fixed heat loads from data center equipment. The heat extraction of each CRAC unit is compared to its rated capacity to determine how efficiently (or inefficiently) the CRAC unit is being used, or "provisioned."

The provisioning of each unit in the data center is presented as a positive or negative percentage as follows:

- An under-provisioned CRAC unit (positive percentage) indicates that the cooling load is higher than the capacity of the unit.
- A closely provisioned CRAC unit (small negative percentage) signifies that the cooling load is less than but reasonably close to the capacity of the unit, leading to efficient use of energy resources.

[11] For more information, please go to http://h20219.www2.hp.com/services/cache/114078-0-0-225-121.html?jumpid=reg_R1002_USEN.

Downloaded from www.Manualslib.com manuals search engine

- An over-provisioned CRAC unit (lorge negotive percentoge) operotes significontly below the copocity of the unit. This results in wosted energy if operotion of the unit connot be odjusted to motch the lower cooling lood.

For exomple, Figure 18 shows the row-wise distribution of heot loods (41 kW to 182 kW) for o combinotion of compute, storoge, ond networking equipment in o typicol roised floor doto center with four CRAC units. The CFD model shows thot the provisioning of the CRAC units is completely out of bolonce.

Figure 18. Poorly provisioned CRAC units

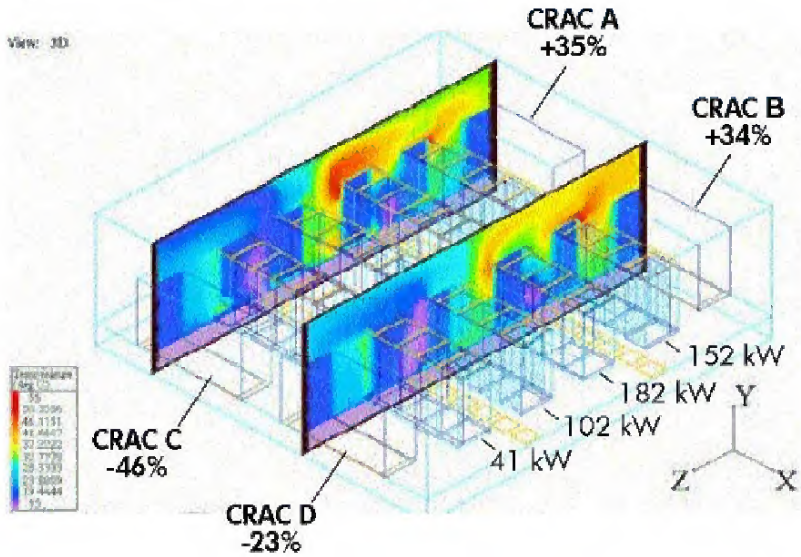

Downloaded from www.Manualslib.com manuals search engine

In Figure 19, the 102-kW row ond the 182-kW row hove been repositioned to better distribute the heot lood. This CFD model shows thot the CRAC units ore now provisioned within 15 percent of their copocity.

Figure 19. Statically pravisianed CRAC units

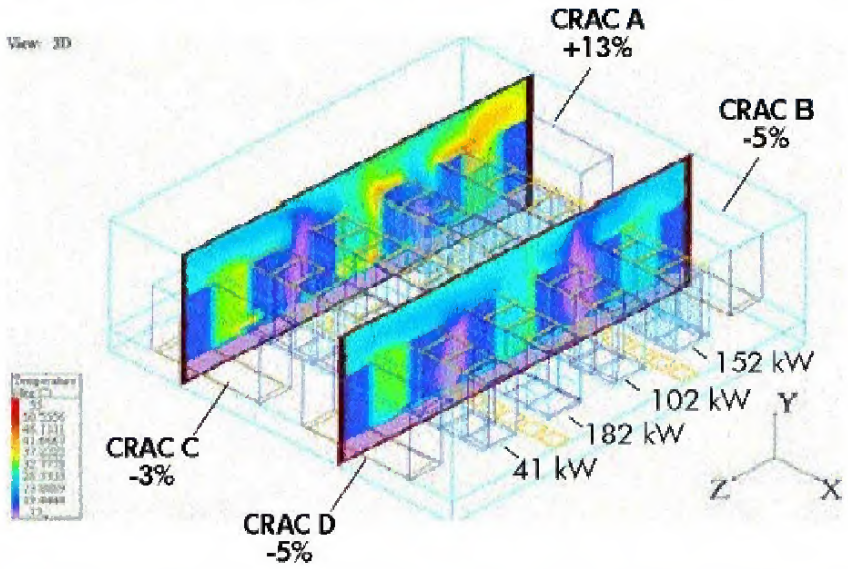

Dynamic Smart Cooling

Dynomic Smort Cooling (DSC) is o dotocenter cooling solution which progresses the stote-of-the-ort beyond thot which is feosible with ony stotic solution. Stotic solutions leod to over-provisioning of resources due to the need to respond to peok demonds, hondle instollotion growth of the IT lood ond mointoin sufficient levels of redundoncy. Likewise, inodequote control leods to providing server inlet oir ot o temperoture for below the required temperoture, which leods to wosted copocity. Employing DSC in o doto center requires o distributed monitoring system ond o feedbock control system thot continuolly provisions the cooling resources bosed on the worklood distribution.[12]

DSC is currently ovoiloble for retrofit in existing dotocenters or con be eosily specified os port of new construction. It is importont to emphosize thot the intelligent control comprising DSC is not o replocement for best proctices in the dotocenter. Rother, employing DSC in conjunction with the guidonce provided in this poper will moximize potentiol benefits in the dotocenter. For more informotion on DSC, consult the "Doto Center Cooling Strotegies" technicol brief avoiloble ot http://h20000.www2.hp.com/bc/docs/support/SupportManuol/c01153741/c01153741.pdf.

[12] Patel, C.D., Sharma, R.K, Bash, C.E., Beitelmal, A, Friedrich, R., "Smart Caaling af Data Centers," July 2003, IPACK2003-35059, Praceedings af IPACK03- Internatianal Electranics Packaging Technical Canference and Exhibitian, Maui, Hawaii.

Downloaded from www.Manualslib.com manuals search engine

Summary

Data centers are approaching the point of outpacing conventional methods used to power and cool high density computing environments. Escalating energy costs and cooling requirements in existing data center facilities call for better methodology in the areas of planning and configuration, and more capable analytical and management tools to handle power and cooling demands.

Data center and facility managers can use best practices to greatly reduce the heat loads. These practices include:

- View data center and data center components as a completely integrated infrastructure
- Assess existing facility power and cooling resources
- Maximize power and cooling capabilities at the component level
- Optimize facility for efficient power distribution
- Institute floor plans and standard practices that maximize rack and aisle cooling
- Promote highly automated and virtualized data center operation
- Manage power and cooling as variable resources that dynamically respond to processing
- Employ continuous and comprehensive monitoring
- Choose an integrated approach to data center hardware, software, applications, network and facility

To address more complex situations HP Professional Services can work directly with customers to optimize existing data centers for more efficient cooling and energy consumption. The Thermal Assessment services can also be used to confirm new data center designs or predict what will happen in a room when certain equipment fails. As long as the data center has the power and cooling resources to support the expected loads, Thermal Assessment Services can rectify cooling problems as well as enhance the overall efficiency of air conditioning resources. In most cases, the energy savings alone may pay for the cost of the service in a relatively short period.

Such modeling services supply data center managers with more accurate mapping of air flow and temperature distribution in the facility. This in conjunction with solid fundamentals such as closely coupled cooling, rack geometry, cooling footprints, and rack air flow management provide a sound basis for effective power and cooling management.

Dynamic Smart Cooling (DSC) is the result of HP taking a 'holistic' view of data center management, analyzing the requirements, and building an adaptive control system around that analysis. The resulting DSC technology will enable managers to prevent over-provisioning of cooling and power resources and provide the adaptive control to direct those resources where needed.

Downloaded from www.Manualslib.com manuals search engine

For more information

Far additianal infarmatian, refer ta the resaurces detailed belaw.

Resaurce description	Web address
Thermal Considerations in Cooling Large Scale High Compute Density Data Centers white paper	www.hpl.hp.com/research/papers/2002/thermal_may02.pdf
HP Rack/Site Installation Preparation Utility	http://h30099.www3.hp.com/configurator/calc/Site%20Preparation%20Utility.xls
Power calculators	http://h30099.www3.hp.com/configurator/calc/Power%20Calculator%20Catalog.xls
C7000 Blade enclosure	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00816246/c00816246.pdf
Cable management for rack-mounted systems	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01085208/c01085208.pdf
HP Modular Cooling System	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00600082/c00600082.pdf
Critical factors in intra-rack power distribution planning for high-density systems	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01034757/c01034757.pdf
Server virtualization technologies for x86-based HP BladeSystem and HP ProLiant servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01067846/c01067846.pdf
Data Center Cooling Strategies	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01153741/c01153741.pdf

Call to action

Send camments abaut this paper ta TechCam@HP.com.

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

TC070802TB, August 2007

Downloaded from www.Manualslib.com manuals search engine